

SYSTEM AND METHOD FOR ANALYZING CAPACITY IN A PLURALITY OF PROCESSING SYSTEMS

Field of the Invention

The invention relates generally to the field of computer systems and more particularly to a system and method for optimizing computer resource usage across a plurality of computer systems.

Background of the Invention

In the capacity planning process, system parameters, desired service levels, and workload predictions are used to determine when the resources of a computer system will be exceeded and are used to assist in identifying cost-effective remedies to resource shortfalls. "Capacity Planning and Performance Modeling: From Mainframes to Client-Server Systems", by Daniel A. Menasce, Virgilio A. Almeida, and Larry W. Dowdy (Prentice Hall, Englewood Cliffs, New Jersey, 1994) discloses approaches to both the predicting and rectifying of computer resource challenges.

Capacity planning for a set of heterogeneous computer systems presents several problems, as set forth below. As a first challenge, it must be recognized that workloads use multiple resources. Therefore, the effect of workload assignment

is not readily predicted or quantified. Second, workload typically grows, and the rate of growth may differ between resources. Third, different computer systems may have different resources, and different resource capacities. These problems can make it difficult to determine how long available resources will last, which computer systems are most at risk for exceeding their resources, how to reallocate resources to alleviate shortages, and how the computer systems will be affected by such reallocations.

007107 27805960
Dan Asit and Dinkar Sitaram, in U.S. Patent No. 5,530,557, entitled "Online Placement of Video Files Determined by a Function of the Bandwidth to Space Ratio of each of the Storage Devices in a Server Environment", (June 25, 1996) teach one solution for maximizing storage utilization for the placement of videos on storage devices taking into account the expected demand for the video. Asit, et al use the bandwidth space ratio (BSR) to place videos on disks. The BSR of a disk is its bandwidth divided by space. The BSR of a video is the expected demand for the video divided by the space required to store it. Demand may be forecast based on historical usage data and, in their invention, a Video Placement Manager places the videos on the disks to match the BSR of the videos with the BSR of the disk.

Additional references which have sought to predict and manage storage capacity include an article and related patent

critical resources have been identified, workload assignment can be more equitably made to improve resource usage.

Brief Description of the Drawings

The invention will now be described in greater detail with specific reference to the appended drawings wherein:

Fig. 1 provides a graph illustrating the mapping of capacity space for two resources of a computer system in units of time in accordance with the present invention;

Fig. 2 illustrates a processing environment for implementing the present invention;

Fig. 3 provides a graph representatively mapping the capacity space based on the life expectancies of resources in a processing environment;

Fig. 4 illustrates a representative critical action timeline for action by the administrative processor of the present invention;

Fig. 5 shows the effect of shifting the workload from system S to system D on the capacity space;

Fig. 6 shows the effect of shifting the workload from system S to system D on a critical action timeline;

Fig. 7 shows a process flow for a workload prioritization procedure for use with the present invention;

A *repository* is a means for storing structured data external to the administrative processor. Data in the repository is saved and accessed in a storage subsystem but is also supported by software, such as relational database software, that provides access to the structure of the data. Database software is not essential for a repository as the content of the repository can be stored in simpler storage objects, frequently called a flat file.

A *processing system* is a computing system that includes all the hardware and software needed to execute computer programs. This includes the central processing unit (CPU) or multiple CPUs, memory, storage and network connectivity as well as the operating system, application software and procedures for managing work on the system.

Workload is the set of identifiable tasks that execute in the processing system and utilize or consume the resources of the system.

A *workload unit* is a subset of the workload that can be associated with some external identifier (e.g., the collection of all tasks executed by an employee user.) Workload units are a collection point for keeping historical records about resource consumption and act as a means to allocate workload to a specific processing system. Workload units may execute anywhere in the processing environment, subject only to resource constraints.

resource needs of this system must be addressed in time CLS. Because all of the resources are represented in units of time, an arbitrary number of dimensions can be collapsed into one in this way.

A system with multiple resources is *balanced with respect to life expectancy* if all of its resources have the same life expectancy. Otherwise, the system is *unbalanced*. The life expectancy of an unbalanced system is the minimum life expectancy over all of its resources. Balanced systems fall on a line in the space drawn from the origin through (n, n, n, \dots, n) , for some constant n , where the size of the tuple is the number of resources or dimensions.

Non-critical resources in a processing system are said to have *slack* beyond the critical resource, $ES_i = LS_i - CLS$. Slack represents available resources that could be reallocated under the present invention.

A capacity space for two resources is shown in Figure 1. The resources are disk storage and CPU capacity. The circle at 110 represents a processing system in which the system is unbalanced since disk storage is expected to suffice for 20 days, and the CPU capacity is expected to suffice for 60 days. The circle at 112 represents a balanced processing system in which both resources are expected to run out in 40 days. In Figure 1, the critical resource for the processing system represented by the circle at 110 is storage; while for the balanced processing

system represented by the circle at 112, both resources are equally critical. Furthermore, in Figure 1, the processing system represented by the circle at 110 has 40 days of slack in its CPU resource.

With reference to Fig. 2, a processing environment where this invention would apply is represented wherein the administrative processor 201 is adapted to implement the inventive process. The administrative processor 201 has access to a data repository 202, such as a relational database, where the data can be saved and from which data can be retrieved. The administrative processor also has access to configuration data about properties of the processing environment 200, comprising processing systems 204a and 204b, and about workload needs and workload usage history data relating to the workload units 205a and 205b that are part of the processing environment.

The objects of interest for this invention are processing systems 204a and 204b that manage workload units. An example of work which is to be divided up into workload units to be managed by the respective processing systems is a sort program run on behalf of a computer user. Each workload unit has a unique identifier, within the processing environment. Workload units consume resources of the processing systems. The resources consumed by a workload unit are recorded by the processing system and the record of this consumption is transferred as workload usage history to the administrative processor 201 and stored in the

repository 202. Each resource has its own unit of measure. These consumption records are identified by the name of the workload unit and the time period of the consumption.

The administrative processor of the present invention utilizes a list of resources in the processing environment, $R=\{R1...Rn\}$; a list of processing systems in the processing environment, $S=\{S1...Sk\}$; for each processing system S_i , its resource capacities $CS_i=\{CR1...CRn\}$, and the workload usage histories stored in repository 202. Once the administrative processor, 201 of Fig. 2, gathers the foregoing information, it constructs a capacity space based on the life expectancy of the resources in the processing environment. Fig. 3 provides an illustration of a two-dimensional capacity space with the CPU life expectancies defined along one axis and the storage life expectancies along the other axis. Under the present invention, a N-dimensional space can be created for N different resources. For purposes of ease of illustration, however, the 2-dimensional space is illustrated and described.

In Fig. 3, a critical resource line is defined at 45°. For a balanced system, such as S_4 , which is plotted at graph point 304, the life expectancies of its resources are equal (80 days) and the graph point necessarily falls on the critical resource line. All of the other systems which are plotted in Fig. 3, S_1 , S_2 , and S_3 , represent unbalanced systems for which one resource has a shorter life expectancy than the other system resource.

When the workload is placed on another so-called destination system D, the shift may decrease the life expectancy of the system D resources, moving from 502 to 512, shown as D'. Note that because the resource capacities and current usage of S and D may differ, the increase in life expectancy on S may not be equal to the decrease in life expectancy on D for a given resource. Figure 6 shows the effect of shifting the workload from source system S to destination system D on the critical resource line (or the critical action timeline), with time 601 shifting back to 611 for source system S while time 602 shifts forward to 612 for destination system D.

A wide variety of existing algorithms can be used to balance resources by shifting workload between them and moving them in capacity space. Without precluding the use of any other such algorithm, a representative prioritization procedure and two representative workload reallocation procedures are set forth in detail below.

Figure 7 shows a flowchart for a prioritization procedure which may be used in conjunction with the present invention. The output of the procedure is a list of systems, or containers, identifying their critical resource, and sorted on the times at which they are expected to expire. When the administrative processor has retrieved the list of systems from the repository, along with the list of resources, the administrative processor begins iterating through the systems S_i , selecting them one at a

time for review and analysis. At step 702 a system is selected followed by selecting a resource at step 703. The life expectancy for each resource is calculated at step 704. This process is repeated until the life expectancy for every resource in a given processing system has been determined. If it is determined at step 705 that there are no other resources in the processing system to be evaluated, the administrative processor then determines the critical resource for the system at 706 as the resource having the minimum life expectancy, and stores the critical resource and its critical time at step 707. Next the administrative processor checks to see if all systems have been evaluated. If not, another system is selected, its resources are analyzed and its critical resource and critical time are stored, as above. If all systems have been evaluated, then the stored critical resource and critical time data are retrieved and sorted by time at step 708. Finally, a timeline is output at 709 which represents the timetable for action described above.

An alternative method, based on a graph of the capacity space, is as follows:

1. Construct the capacity space
2. Plot the critical resource line
3. Plot the points in the capacity space
4. Project the points onto the critical resource line at the life expectancy of their critical resource

